CLAIMS

What is claimed is:

1. A method for replicating data from a master server to a slave server over a network, the method comprising the steps of:

sending a packet of information from the master server to the slave server, the information relating to a change in the data stored on the master server and containing a version number for the present state of the data;

allowing the slave server to determine whether the data on the slave server has been updated to correspond to the version number contained in the packet; and

requesting a delta be sent from the master server to the slave server if the data on the slave server does not correspond to the version number contained in the packet, the delta containing information needed to update the slave server.

2. A method according to claim 1, further comprising:

storing an original copy of the data on the master server.

3. A method according to claim 1, further comprising:

persistently caching the data on a local disk for each slave server.

4. A method according to claim 1, further comprising:

determining a unique version number for the current state of the data on the master server if the data has changed.

5. A method for replicating data from a master server to a slave server over a network, the method comprising the steps of:

sending a version number from the master server to the slave server, the version number relating to the present state of the data stored on the master server;

allowing the slave server to determine whether the slave server has been updated to reflect the present state of the data corresponding to the version number sent from the master server; and

requesting a delta be sent from the master server to the slave server if the slave server does not correspond to the version number sent by the master, the delta containing information needed to update the slave server.

6. A method according to claim 5, further comprising:

sending the delta from the master server to the slave server.

7. A method according to claim 5, further comprising:

committing the delta to the slave server.

8. A method according to claim 5, further comprising:

updating the version number of the slave server after committing the delta.

9. A method according to claim 5, further comprising:

periodically sending the version number from the master server to a slave server.

10. A method according to claim 5, further comprising:

sending the version number to a slave server until the slave server acknowledges receipt of the version number.

11. A method according to claim 5, further comprising:

including data with the version number that is necessary to update a slave server.

12. A method according to claim 11, further comprising:

committing the data necessary to update the slave server as soon as it is received.

13. A method according to claim 5, further comprising:

determining the scope of the delta before sending it from the master server.

14. A method for replicating data over a network including a master server and at least one slave server, the method comprising the steps of:

sending a packet of information from a master server to each slave server on the network, the information relating to a change in the data stored on the master server and containing a current version number for the present state of the data, the information further relating to previous changes in the data and a version number for each previous change;

allowing each slave server to determine whether the slave server has been updated to correspond to the current version number;

allowing each slave server to commit the information if the slave server has not missed a previous change; and

allowing each slave server having missed a previous change to request that previous change be sent from the master server to the slave server before the slave server commits the packet of information.

15. A method according to claim 14, further comprising:

committing the packet of information to a slave server.

16. A method according to claim 14, further comprising:

aborting the commit of the packet of information if a slave server cannot commit the update.

17. A method according to claim 14, further comprising:

determining the scope of the delta before sending it from the master server.

18. A method according to claim 14, further comprising:

including the scope of each the previous changes in the delta.

19. A method for replicating data over a network including a master server and at least one slave server, the method comprising the steps of:

sending a packet of information from a master server to each slave server on the network, the information relating to a change in the data stored on the master server and containing a prior version number for the prior state and a new version number for the new state of the data, the information further relating to previous changes in the data and a previous version number

for each previous change;

allowing each slave server to determine whether the data on the slave server corresponds to the prior version number contained in the packet;

allowing each slave server to commit the packet of information if the data on the slave server corresponds to the prior version number contained in the packet, the commit also updating the version of the slave server to the new version number; and

allowing each slave server not corresponding to the prior version number to request that a delta be sent from the master server containing the information necessary to update the slave to the prior version number before the slave server commits the packet of information.

20. A method for replicating data over a network including a master server and at least one slave server, the method comprising the steps of:

sending a packet of information from a master server to each slave server on the network, the information relating to a change in the data stored on the master server and containing a version number for the prior state and a version number for the new state of the data, the

information further relating to previous changes in the data and a version number for each previous change;

allowing each slave server to determine whether the data on the slave server corresponds to the prior version number contained in the packet;

allowing each slave server to commit the packet of information if the data on the slave server corresponds to the prior version number contained in the packet, the commit also updating the version of the slave server to the new version number; and

allowing each slave server not corresponding to the prior version number to request that a delta be sent from the master server containing the information necessary to update the slave to the new version number.

21. A method for replicating data from a master server to at least one slave server over a network, the method comprising the steps of:

sending a packet of information from the master server to a slave server, the information relating to a change in the data stored on the master server and containing a version number for the present state of the data;

receiving the packet of information to a slave server;

allowing the slave server to determine whether the slave server has been updated to correspond to the version number contained in the packet, and to further determine whether the slave server can process the packet of information if needed to update to correspond to the version number contained in the packet;

sending a signal from the slave server to the master server, the signal indicating whether the slave server needs to be updated and whether the slave server can process the update; and

sending a response signal from the master server to the slave server indicating whether the slave server should commit to the information contained in the packet; and

committing the packet of information to the slave server if so indicated by the response signal.

22. A method according to claim 21, further comprising:

determining whether each of the at least one slave server can commit the data.

23. A method according to claim 21, further comprising:

determining whether each of the at least one slave server has sent a response back to the master server.

24. A method according to claim 21, further comprising:

determining whether any of the at least one slave server can commit the data.

25. A method according to claim 21, further comprising:

committing the data only if each of the at least one slave server can process the commit.

26. A method according to claim 21, further comprising:

aborting the data only if any of the at least one slave server cannot process the commit.

27. A method according to claim 21, further comprising:

committing the data to those slaves that are able to process the commit.

28. A method according to claim 21, further comprising:

multicasting the update to any of the at least one slave server that were not able to process the commit.

29. A method according to claim 21, further comprising:

heartbeating the new version number to any of the at least one slave server that were not able to process the commit.

30. A method according to claim 21, further comprising:

requesting a delta be sent to a slave server that was not able to process the commit.

31. A method for replicating data over a network, the method comprising the steps of:

(a) determining whether the replication should be accomplished in a one or two phase method;

(b) sending replication information determined to he accomplished in a one phase method by:

sending a packet of information from the master server to the slave server, the information relating to a change in the data stored on the master server and containing a version number for the present state of the data;

receiving the packet of information to a slave server;

allowing the slave server to determine whether the data on the slave server has been updated to correspond to the version number; and

requesting a delta be sent from the master server to the

slave server if the slave server does not correspond to the version number, the delta containing information needed to update the slave server;

(c) sending replication information determined to he accomplished in a two phase method by:

sending a packet of information from the master server to the slave server, the information relating to a change in the data stored on the master server and containing a version number for the present state of the data;

allowing the slave server to determine whether the slave server has been updated to correspond to the version number, and to further determine whether the slave server can process the packet of information;

sending a signal from the slave server to the master server indicating whether the slave server needs to be updated and whether the slave server can process the packet of information;

sending a response signal from the master server to the slave server indicating whether the slave server should commit to the packet of information; and

committing the packet of information to the slave server if so indicated by the response signal.

32. A method for replicating data over a network, the method comprising the steps of:

(a) determining whether replication should be accomplished in a one or two phase method;

(b) sending data to be replicated in a one phase method by:

sending a version number for the current state of the data from a master server to a slave server;

requesting a delta be sent from the master server to the slave server if the data on the slave server does not correspond to the version number; and

(c) sending data to be replicated in a two phase method by:

sending a packet of information from the master server to a slave server;

determining whether the slave server can process the packet of information; and

committing the packet of information to the slave server if the slave server can process the packet of information.

33. A method for replicating data from a master to a plurality of slaves on a network, the method comprising the steps of:

(a) determining whether replication should be accomplished in a one or two phase method;

(b) sending data to be replicated in a one phase method by:

sending a version number for the current state of the data from the master to each slave; and

requesting a delta be sent from the master to each slave containing data that does not correspond to the version number;

(c) sending data to be replicated in a two phase method by:

sending a packet of information from the master to each slave; and

committing the packet of information to the slaves if each of the plurality of slaves can process the packet of information.

34. A method for replicating data from a master to a plurality of slaves on a network using one and two phase methods, the method comprising the steps of:

(a) sending data to be replicated in a one phase method by sending a version number for the current state of the data from the master to each slave so that each slave may request a delta to be sent from the master to the slave to update the data on the slave; and

(b) sending data to be replicated in a two phase method by sending a packet of information from the master to each slave, the packet of information to be committed by each slave if every

slave is able to commit the packet of information.

35. A method for replicating data on a clustered network using one and two phase methods, each network cluster containing a cluster master and at least one cluster slave, the method comprising the steps of:

(a) sending data to be replicated in a one phase method by sending a version number for the current state of the data from a first cluster master to all other cluster masters so the other cluster masters may each request a delta; and

(b) sending data to be replicated in a two phase method by sending a packet of information from the first cluster master to each other cluster master, the packet of information to be committed by the other cluster masters if the other cluster masters are able to commit the packet of information.

36. A method according to claim 35, further comprising:

sending the data from each cluster master to each cluster slave in the cluster with that cluster master by a one-phase method.

37. A method according to claim 10, further comprising:

sending the data from each cluster master to each cluster slave in the cluster with that cluster master by a two-phase method.

38. A computer-readable medium, comprising:

(a) means for sending a packet of information from a master server to each slave server on the network, the information relating to a change in the data stored on the master server and containing a current version number for the present state of the data, the information further relating to previous changes in the data and a version number for each previous change;

(b) means for allowing each slave server to determine whether the slave server has been updated to correspond to the current version number;

(c) means for allowing each slave server to commit the information if the slave server has not missed a previous change; and

(d) means for allowing each slave server having missed a previous change to request that previous change be sent from the master server to the slave server before the slave server commits the packet of information.

39. A computer program product for execution by a server computer for replicating data over a network, comprising:

(a) computer code for sending a packet of information from a master server to each slave server on the network, the information relating to a change in the data stored on the master server and containing a current version number for the present state

of the data, the information further relating to previous changes in the data and a version number for each previous change;

(b) computer code for allowing each slave server to determine whether the slave server has been updated to correspond to the current version number;

(c) computer code for allowing each slave server to commit the information if the slave server has not missed a previous change; and

(d) computer code for allowing each slave server having missed a previous change to request that previous change be sent from the master server to the slave server before the slave server commits the packet of information.

40. A system for replicating data over a network, comprising:

(a) means for sending a packet of information from a master server to each slave server on the network, the information relating to a change in the data stored on the master server and containing a current version number for the present state of the data, the information further relating to previous changes in the data and a version number for each previous change;

(b) means for allowing each slave server to determine whether the

slave server has been updated to correspond to the current version number;

(c) means for allowing each slave server to commit the information if the slave server has not missed a previous change; and

(d) means for allowing each slave server having missed a previous change to request that previous change be sent from the master server to the slave server before the slave server commits the packet of information.

41. A computer system comprising:

a processor;

object code executed by said processor, said object code configured to:

(a) send a packet of information from a master server to each slave server on the network, the information relating to a change in the data stored on the master server and containing a current version number for the present state of the data, the information further relating to previous changes in the data and a version number for each previous change;

(b) allow each slave server to determine whether the slave server has been updated to correspond to the current version number;

(c) allow each slave server to commit the information if the slave server has not missed a previous change; and

(d) allow each slave server having missed a previous change to request that previous change be sent from the master server to the slave server before the slave server commits the packet of information.